

REMARKS

The Examiner's action dated April 21, 2004, has been received, and its contents carefully noted.

This application contains claims 1, 3 and 5-39. Claims 14-39 were withdrawn, with traverse, in response to a restriction requirement. Claims 1, 3 and 5-13 were rejected in the Official Action of April 21. Reconsideration is respectfully requested.

Applicant thanks Examiner LeRoux for the courtesy of an interview which was granted to Applicant's representative, Sanford T. Colb (Reg. No. 26,856). The interview was held in the USPTO on June 29, 2004. In the interview, Mr. Colb discussed the distinction of claim 1 over the combination of references cited by the Examiner (Lasser et al. and Lawthers), and pointed out that the Examiner's reasons for rejection of claim 1 over these references in the present Official Action were inconsistent with statements made in the preceding Official Action. The Examiner agreed to reconsider the basis of the rejection.

Claims 1, 3, 5 and 11-13 were rejected under 35 U.S.C. 103(a) over Lasser et al. (U.S. Patent 5,897,638) in view of

Lawthers ("The Data Management Applications Programming Interface"). Applicant respectfully traverses this rejection.

Claim 1 (as amended in response to the previous Official Action) is drawn to a method for managing data storage, in which a session of a data management (DM) application is initiated on a session node in a cluster, using a data management application programming interface (DMAPI) of a parallel file system in the cluster. A user application runs on a source node in the cluster. The user application running on the source node submits a request to the parallel file system to perform a file operation on one of the files in the data storage. In response to this request, a notification of a DM event is sent to the session node for processing by the DM application. The event is processed at the session node in order to obtain a data management access right, which is then used in performing the file operation.

Lasser describes a parallel virtual file system, which is used to manage partitioned files. Lawthers describes a conventional, single-node DMAPI, and makes no reference to parallel file systems. As noted by the Examiner (page 6, fourth paragraph, in the previous Official Action, in reference to claim 4), neither Lasser nor Lawthers discloses

the use of source and session nodes as recited in claim 1 as  
it is currently presented:

"The combination of Lasser '638 and Lawthers fails to disclose wherein receiving the request comprises receiving an invocation of a file operation submitted to the parallel file system by a user application on a source node, and sending a notification of a DM event to the session node responsive to the request, and wherein obtaining the right at the session node comprises processing the event at the session node subject to the access right [paragraph 43]."

Applicant believes this assessment of the prior art to be correct, and to provide sufficient reason for allowability of the present claims.

In view of the above statement by the Examiner in the previous Official Action, Applicant was surprised and puzzled by the fact that in the present Official Action, amended claim 1, which recites source and session nodes in exactly the manner of original claim 4, was rejected over Lasser and Lawthers. This finding by the Examiner is in direct contradiction to the Examiner's statement in the previous official action, and is not supported by the cited references.

In view of the absence of key limitations of claim 1 from the cited references, Applicant believes that the Examiner has failed to make a *prima facie* case of obviousness, as required by MPEP 2143.03:

"All Claim Limitations Must Be Taught or Suggested.

To establish *prima facie* obviousness of a claimed invention, all the claim limitations must be taught or suggested by the prior art. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA 1974). 'All words in a claim must be considered in judging the patentability of that claim against the prior art.' *In re Wilson*, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970)."

Specifically, regarding Lasser, the Examiner states (on page 10 of the present official action) that Lasser's multifile server reads on the source node recited in claim 1, while the computers containing the control and data trees read on the session node, as disclosed by Lasser in col. 3, line 64 - col. 4, line 13. The cited passage describes a "distributed service" usage pattern of Lasser's system, which is shown by Lasser in Fig. 6, and is described in greater detail in col. 20, lines 1-24. In Fig. 6, it will be observed that program instances 7 are associated with data trees 1, which presumably led the Examiner to identify these program instances with the

"data management application" running on a session node as recited in the first step of claim 1.

There is no indication in Fig. 6 or in the text of Lasser's disclosure to suggest that a user application (or any sort of application) runs on multifile servers 4, as required by the second step of claim 1. In fact, Lasser does not even relate to servers 4 as separate computers, but rather as sets of multifile subroutines 4 (col. 5. lines 14-19), which may run on any or all of the computers using the parallel file system. Fig. 6 shows only "multifile stubs" 9, which do not run on servers 4 but merely "contact the appropriate servers to perform multifile operations" (col. 20, lines 11-12, emphasis added). The multifile servers are just that: file servers, which perform the basic function of providing access to files in the file system. They run no applications of their own.

Furthermore, the Examiner's association of the session and source nodes with certain elements of Lasser's system does not accord with the specific grounds of rejection of claim 1 stated by the Examiner on page 3 of the present official action:

a) The Examiner identifies the step of "creating a session of a data management application on a session node..."

with col. 15, lines 14-29, in Lasser. This passage, however, refers to multifile subroutines, i.e., to Lasser's multifile servers, and to the file operations that they perform. As noted above, however, elsewhere in the official action the Examiner associates the session node not with the servers, but with "computers containing the control and data trees." There is an internal contradiction in the Examiner's reasoning here. Furthermore, there is no mention or suggestion in this passage of anything that could be considered a session, let alone a session of a data management application as recited in claim 1.

b) The Examiner identifies the step of "receiving a request submitted... by a user application on a source node..." with subroutines 4 in Lasser's Fig. 1. But these are the same subroutines as are described in col. 15, lines 14-29, which the Examiner also associated with the first step of claim 1. This association appears to mean that the Examiner considers the data management application and the user application recited in claim 1 to be identical. Claim 1, however, makes a clear distinction between the two types of applications, which may run simultaneously in the cluster. (This is the whole point of the claimed method - to resolve difficulties that arise when user applications and data management applications run on different nodes in the same cluster using a parallel

file system.) Furthermore, although the cited passage in Lasser mentions file operations, it makes no mention of any sort of request by a user application to perform such an operation.

c) The Examiner identifies the step of "sending a notification of a DM event" with Lasser's RPC, and then identifies "obtaining a data management access right" with Lasser's transactional lock. Claim 1, however, requires that the access right be obtained "by processing the event at the session node." By contrast, Lasser's lock is obtained even before the RPC is issued (col. 9, lines 16-20), i.e., obtaining the lock is a prerequisite to issuing the RPC. Thus, even if the RPC were somehow to be considered a "notification of a DM event," Lasser does not meet the requirement that this event be processed in order to obtain the access right, as recited in claim 1.

Thus, Applicant believes that the Examiner was correct in his initial finding that Lasser does not disclose source and session nodes as required by claim 1. On the other hand, even if certain elements of Lasser's system were considered to read on source and session nodes, as the Examiner maintains in the present Official Action, the operation of these elements is very far from meeting the requirements of claim 1.

In the present Official Action, the Examiner goes on to state that Lawthers also discloses source and session nodes, and quotes the last paragraph on the first page of Lawthers in support. Applicant respectfully points out that the cited passage makes no mention of either source or session nodes, nor has the Examiner indicated which element in this passage he believes to be the source node, and which is the session node.

The Examiner continues (on pages 11-12 of the present official action) by rejecting Applicant's arguments against Worfolk (cited against claim 4 in the first official action). As Worfolk is not cited in the present official action, these arguments are now moot, but Applicant continues to maintain their correctness.

Thus, for the reasons stated above, Applicant respectfully submits that claim 1 is patentable over the cited art. In view of the patentability of claim 1, claims 3, 5 and 11-13, which depend from claim 1, are believed to be patentable, as well.

Claims 6-10 were rejected under 35 U.S.C. 103(a) over Lasser in view of Lawthers, and further in view of Wecker (U.S. Patent 6,289,464) and/or Molloy (US 6,625,601). Applicant respectfully traverses this rejection. In view of



the patentability of claim 1, claims 6-10, which depend from claim 1, are believed to be patentable, as well.

Furthermore, Applicant believes that the dependent claims in this application recite independently-patentable subject matter. With respect to claim 6, for example, this claim adds the limitation that a data management lock acquired from the DMAPI is held over a sequence of multiple kernel calls in the parallel file system. Applicant conceded in response to the first official action in this case that Wecker describes making a kernel call to lock certain driver code and buffers (col. 19, lines 9-11). But claim 6 requires not only a kernel call and a lock, but also that the data management lock be held over a sequence of multiple kernel calls. Wecker deals only with single kernel calls, and makes no mention of a multi-call lock as required by this claim. In the present Official Action, the Examiner simply restated his previous rejection of claim 6, and ignored this key added limitation. As this limitation is neither taught nor suggested by the prior art, Applicant believes claim 6 to be independently patentable.

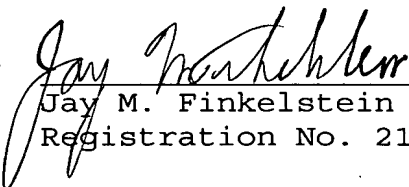
Applicant believes the amendments and remarks presented hereinabove to be fully responsive to all of the grounds of rejection raised by the Examiner. In view of these amendments

Appln. No. 09/887,549  
Reply to Office Action of April 21, 2004

and remarks, applicant respectfully submits that all of the  
claims in the present application are in order for allowance.  
Notice to this effect is hereby requested.

Respectfully submitted,

BROWDY AND NEIMARK, P.L.L.C.  
Attorneys for Applicant

By   
Jay M. Finkelstein  
Registration No. 21,082

JMF:dtb  
Telephone No.: (202) 628-5197  
Facsimile No.: (202) 737-3528  
G:\BN\C\colb\Loy4\PTO\Response 15JUL04.doc